Replayer User Manual

Alistair Morrison morrisaj@dcs.gla.ac.uk

Replayer is a software tool designed primarily to aid in the evaluation of mobile applications. It allows an analyst to combine both videos and system log data in one synchronised system. An activity log recorded on a mobile device will reliably note when an action was performed; video data will show environmental or social factors that might influence a user's behaviour. An analyst can view all the recorded information simultaneously, with several video streams and activity logs from multiple users.

Replayer architecture

Replayer uses an extensible distributed component architecture. Each component of the system runs in its own process, communicating via a TCP interface with the server, and is literally a separate program. The system is based on a client-server paradigm to provide data from a centrally managed database to each of several visualisation components. Replayer's central repository of data is an SQL capable database. Data are stored in a number of tables, with data pushed to each of the visualisation components on demand.

Replayer server

Replayer is designed as a distributed system that can run multiple processes across several machines. One server should be run to manage the database and communication between the various visualisation components. On launching the Server, a view is displayed that summarises the data currently loaded into the system and the number of machines currently visualising the data.

There are two means of loading data into Replayer. An existing project can be loaded from a jar file by selecting from the menubar *File : Open project*. To create a new database, select *File : New project* then system log files can be read into the database by

running a parser (*File : import logfile*). The *File* menu also contains a command to save to a jar.

Parsing system log files

The 'log import' dialog shows all the available parsers and the data format that each expects.

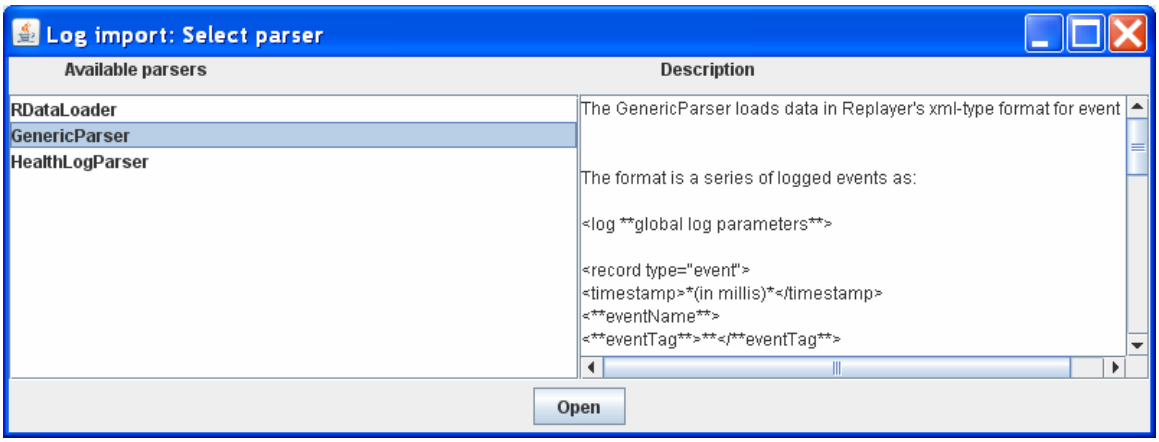

The easiest way to load system log data into Replayer is to generate the logs in a format that is suitable for the existing parsers, as described in the 'log import' dialog.

Replayer Control Unit

The Control Unit can be launched by clicking the button at the bottom of the Server's frame. The purpose of the Control Unit is to allow the user to manage the visualisation components. It is intended that every machine running Replayer should have a Control Unit open (although, as mentioned above, only one machine should run a Server) and that it should be set to always be the topmost window (from the menu *Controls : Always on top*) to help to manage the various other windows that can be open when using multiple coordinated visualisations.

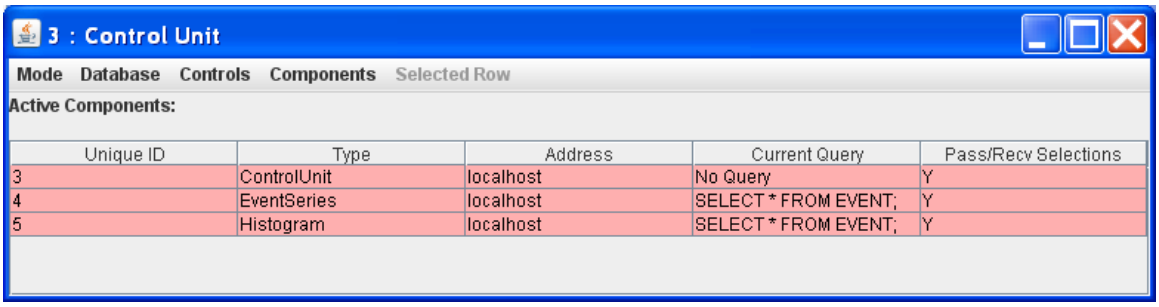
3 : Control Unit
Mode Database Controls Components Selected Row
Active Components:
Unique ID | Type | Address | Current Query | Pass/Recv Selections
3 | ControlUnit | localhost | No Query | Y
4 | EventSeries | localhost | SELECT * FROM EVENT; | Y
5 | Histogram | localhost | SELECT * FROM EVENT; | Y

The main view of the Control Unit is the table of currently running components, showing on which machine it is running and the data it is visualising. New components are opened through the menu *Components : Open component*. This launches the

visualisation process in a new frame and adds a new row to the active components table. Clicking on a row of this table with the middle mouse button minimises/restores the component (and greys out the line in the table if the component is not visible).

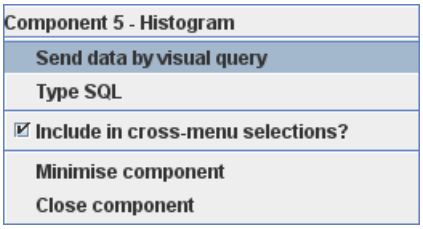

Right-clicking a row of the table opens a pop-up menu from which a number of commands can be made. Data is loaded into components by querying the database via SQL commands. These can be typed in directly, or auto-generated via a visual interface by selecting *Send data by visual query*.

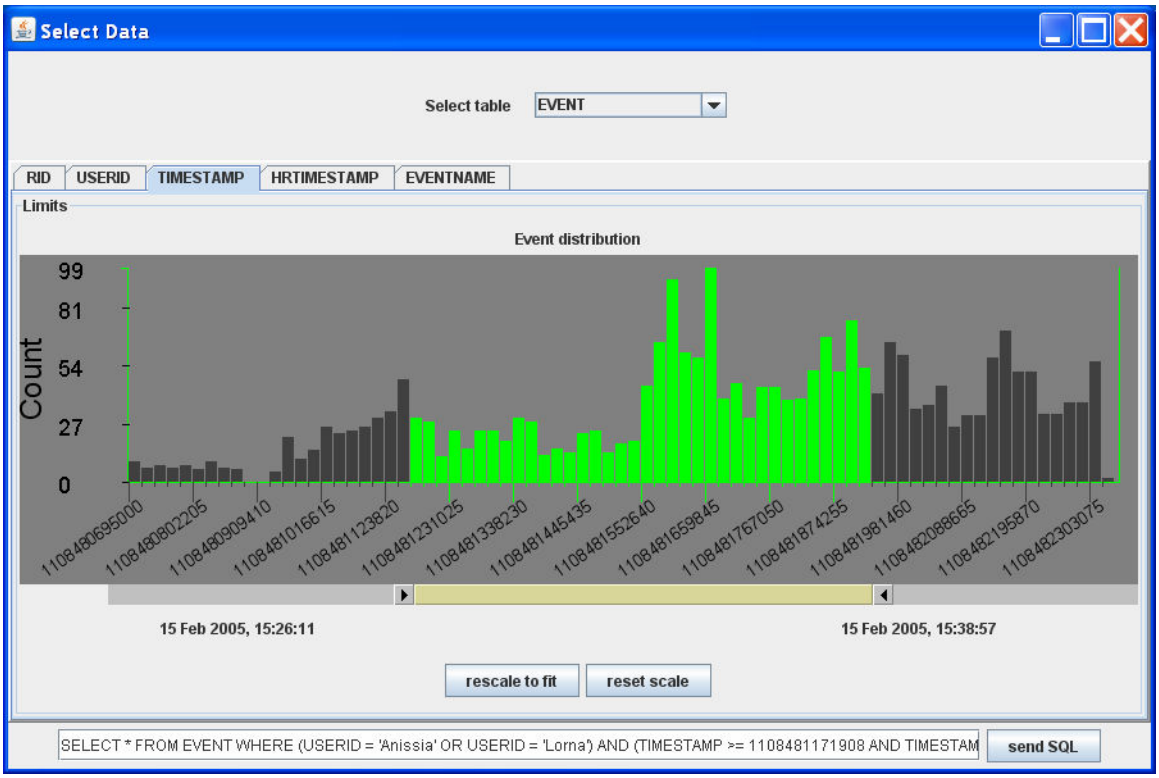

The visual query tool provides a graphical interface to the contents of the database tables and can allow the user to filter data without having to directly compose SQL queries. All the tables in the database are shown in a drop-down menu at the top of the window. On selecting a table, each of its columns are shown as tabbed panes below. If the column contains nominal data (eg Strings) then each unique value is shown as a checkbox; checking any of these boxes means that the returned data will be filtered to contain only the rows with the selected values. Conversely, if the column contains only numerical data then a distribution of values is shown, as in the screenshot above. These can be

filtered to only return specific values using the double-ended slider below the distribution.

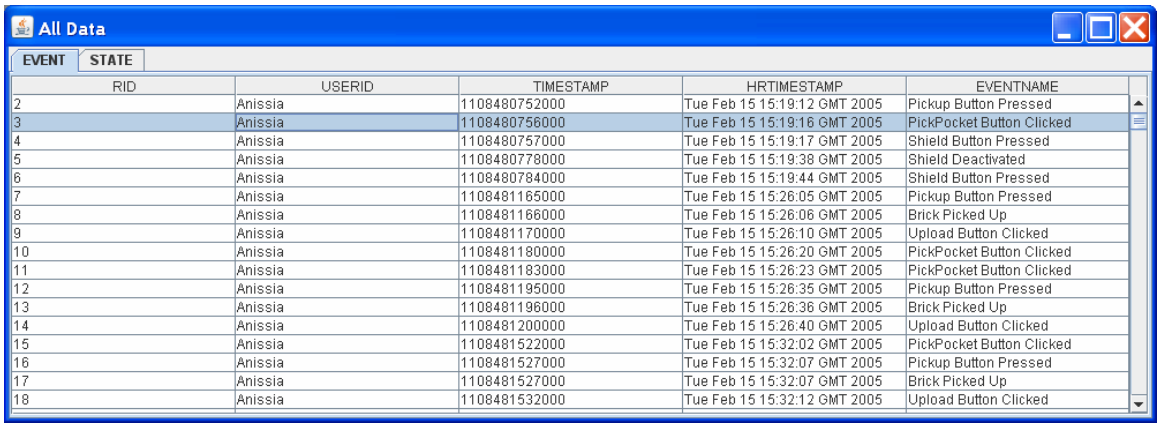

RID	USERID	TIMESTAMP	HRTIMESTAMP	EVENTNAME
2	Anissia	1108480752000	Tue Feb 15 15:19:12 GMT 2005	Pickup Button Pressed
3	Anissia	1108480756000	Tue Feb 15 15:19:16 GMT 2005	PickPocket Button Clicked
4	Anissia	1108480757000	Tue Feb 15 15:19:17 GMT 2005	Shield Button Pressed
5	Anissia	1108480778000	Tue Feb 15 15:19:38 GMT 2005	Shield Deactivated
6	Anissia	1108480784000	Tue Feb 15 15:19:44 GMT 2005	Shield Button Pressed
7	Anissia	1108481165000	Tue Feb 15 15:26:05 GMT 2005	Pickup Button Pressed
8	Anissia	1108481166000	Tue Feb 15 15:26:06 GMT 2005	Brick Picked Up
9	Anissia	1108481170000	Tue Feb 15 15:26:10 GMT 2005	Upload Button Clicked
10	Anissia	1108481180000	Tue Feb 15 15:26:20 GMT 2005	PickPocket Button Clicked
11	Anissia	1108481183000	Tue Feb 15 15:26:23 GMT 2005	PickPocket Button Clicked
12	Anissia	1108481195000	Tue Feb 15 15:26:35 GMT 2005	Pickup Button Pressed
13	Anissia	1108481196000	Tue Feb 15 15:26:36 GMT 2005	Brick Picked Up
14	Anissia	1108481200000	Tue Feb 15 15:26:40 GMT 2005	Upload Button Clicked
15	Anissia	1108481522000	Tue Feb 15 15:32:02 GMT 2005	PickPocket Button Clicked
16	Anissia	1108481527000	Tue Feb 15 15:32:07 GMT 2005	Pickup Button Pressed
17	Anissia	1108481527000	Tue Feb 15 15:32:07 GMT 2005	Brick Picked Up
18	Anissia	1108481532000	Tue Feb 15 15:32:12 GMT 2005	Upload Button Clicked

Another Control Unit menu command (*Database : show full data*) allows the full database tables to be viewed, as shown above.

Playback / Overview modes

Replayer has two main modes operation, as set in the Control Unit's *Mode* menu. Each visualisation component can be in one of the two modes, and will display data differently depending on which mode is set. Playback mode is intended to augment video playback. Several synchronised videos can play at once, and the other components update as the video plays to show information relevant to the current point in the recorded footage. For example, map components give a constantly updating display of participants' current locations, or other visualisations show information on current system state. The second mode is 'overview mode', where the components show visualisations and statistical summaries of all recorded data. In this mode, all the recorded data may be analysed in the same view, rather than only that at the point of video playback. For example, a map tool can display a 'trail' that draws a line following the route of the participant's location throughout the trial. Components are linked, so that selections made in one are reflected in all the others, eg selecting a particular user's column on a histogram would then highlight that user's locations on a map. This applies to video data as well, so that selections and filters (eg by users, locations or event type) can allow the user to find the periods of footage of interest.

The selection process works as follows. A selection is made in a visualisation and the specific rows of database tables selected are sent to the server as a list of comma-separated unique values (the RIDs, one of which is assigned to every table row on data entry). The server then sends out these same RIDs to all other running visualisations, which then highlight the selection (by whatever means they wish). If all visualisations are loaded with the same data, then the highlighting is simply done on the same rows as the selection. If different visualisations have data from different tables, the selections are based on timestamps: the highlighting shows objects recorded in different tables at

similar times. The period over which the selections are linked can be set using *Database : Set selection association timeframe* on the Control Unit.

Any visualisation can 'opt out' of sending or receiving selections by right-clicking on the component on the Control Unit's table and unchecking the checkbox marked *Include in cross-table selections*.

Time Series component

The Time Series component draws each numerical data variable provided as input as a line on the plot, with colours related to variable names by a key on the right. Axes are labelled automatically, and re-labelled if the Time Series module window is resized.

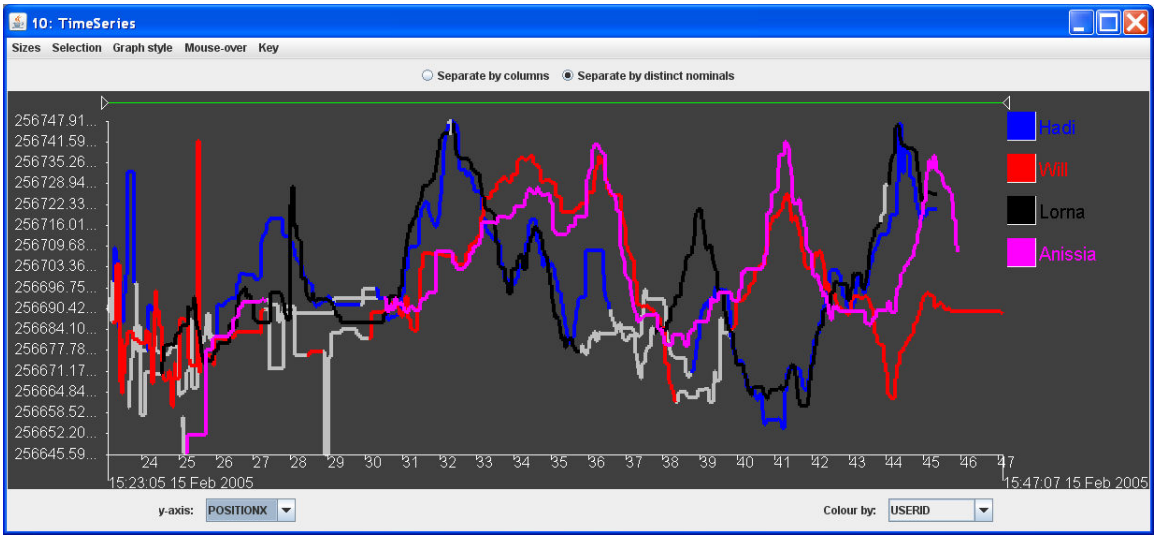

The Time Series module can be used in combination with other visualisations, to allow selection and linking. A double-ended slider is provided at the top of the module, with which a specific section of the series can be selected. If the data are also being viewed in another connected component, the system will highlight the corresponding time measurements. Similarly, marking time periods in other views will alter the Time Series display to reflect the selection. In the screenshot, the Time Series is reflecting selections made in another component. The IDs of the objects contained in the selection have been passed to the Time Series so that the corresponding areas of the display could be highlighted. The Time Series variables are greyed out at the deselected periods.

Placing the mouse over any point of a line on the plot colours the whole line in bright green to help the user to trace it from start to finish. This behaviour can be disabled from the *Mouse*-over menu, to avoid slow-down if a very large data set is being visualised. Other menu commands allow the key to be launched in a separate frame if desired, or to change graph styles. Right-clicking the Time Series displays a pop-up menu to access commands such as resizing axis text and line thickness or to export the image as a PNG file.

Event Series component

The time series module shows system state over time. The Event Series is provided to visualise discrete events that happen occasionally throughout system use, such as user interactions like button clicks. The Event Series has the same time-based x-axis as the time series, but visualises these discrete events as icons.

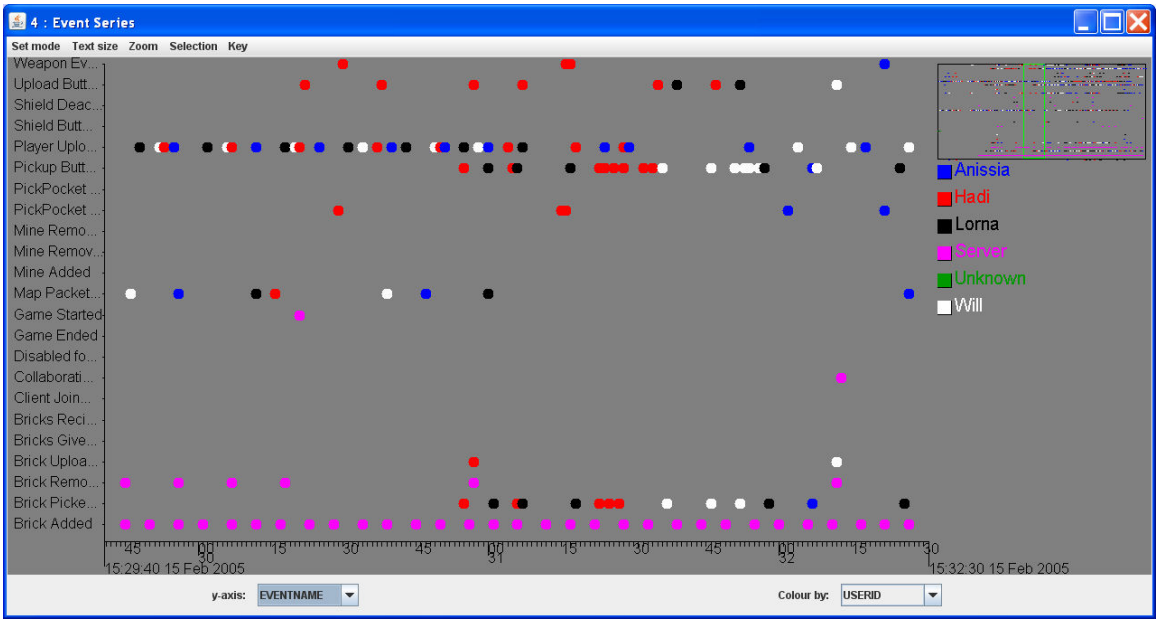

The tool is useful in providing an immediate overview of a whole trial, whilst allowing users to zoom to particular periods of interest. A context window in the top right of the tool shows all of the data currently loaded into the tool, with a green window illustrating where the current focus fits into the overall context. The green window can be dragged around in the context view to immediately jump to another area of interest. Drawing a rectangle in the context view with the middle mouse button sets that region in the main view.

A drop-down list is provided at the bottom-left of the tool to allow users to select the input dimension with which to plot the data in the y-axis. If a numerical dimension is selected, the y-axis will be scaled and each event will be drawn at the appropriate height in the frame. On the other hand, if a nominal dimension is selected, the number of unique nominal values on that dimension is calculated, these are spaced out evenly on the y-axis and objects are placed accordingly. At the bottom-right of the tool is another drop-down list, to determine the input dimension with which to colour objects. A key is provided, as in the time series. It can be used to filter data by clicking one of the boxes to remove that nominal from the plot.

Selections can be made in the Event Series by dragging a box around the objects of choice. Non-selected items are then greyed out.

Histogram

The previously described components afforded temporal-based distributions, and judged events and states on properties at a given time. Replayer also contains a histogram tool, which provides a means of assessing data by distribution. Rather than showing individual events, a summary is given of the aggregated measurements.

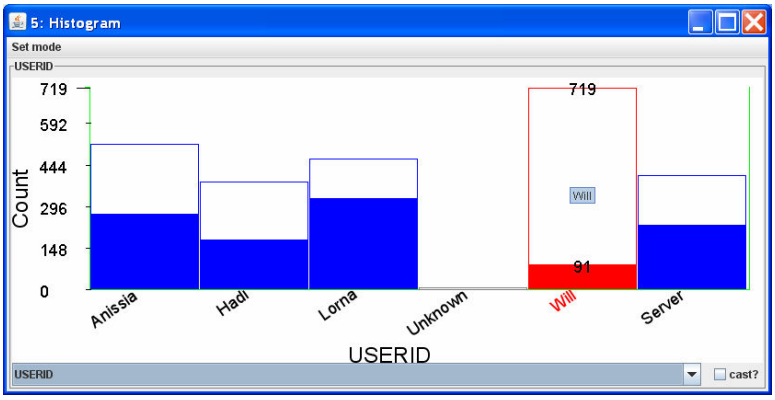

Like the Event Series, the histogram tool has a dropdown list from which an input dimension can be selected. The histogram will then display the distribution of values recorded on that dimension. Should the selected attribute be based on nominal data, a separate bar in the histogram is created for each unique value. This is the case in the figure, which illustrates the activity of each distinct user. If the selected dimension contains numerical data, the data will be bucketed by value, as shown in the visual query generator above. Tool tips are used to present the *x*-axis labels, with the appropriate nominal being displayed when the cursor is moved over a bar. This action also shades the bar red and highlights its height in red for easy value comparisons.

As with the other components, support is provided for brushing between views. Selections can be made by clicking on individual bars, which will instruct the histogram to send other components the IDs of the objects represented therein. The histogram visualises received selections by shading certain amounts of each bar. A bar's colour is filled in proportion to the number its constituent objects that are selected. In the screenshot, most of Lorna's events have been selected, whereas relatively few are selected for Will.

In 'playback mode', the histogram can be set to show certain aspects of system state. During playback, the values will change to give readings in synchronisation with the video. Two combo boxes are provided to select the nominal values used as the *x*-axis and the variable to graph (in the screenshot below, user ID and wifi signal strength respectively). As playback progresses, values at the top of the component show the time at which the video was recorded.

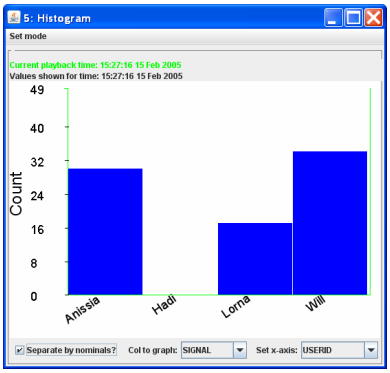

Video component

Replayer plays videos with an Apple OSX-only component using QuickTime media player. The area of the slider's track coloured red shows that area where media clips are available while the green blocks are those areas of time highlighted by making a selection in another tool. The thumb of the slider points to the current point in time being displayed on the videos.

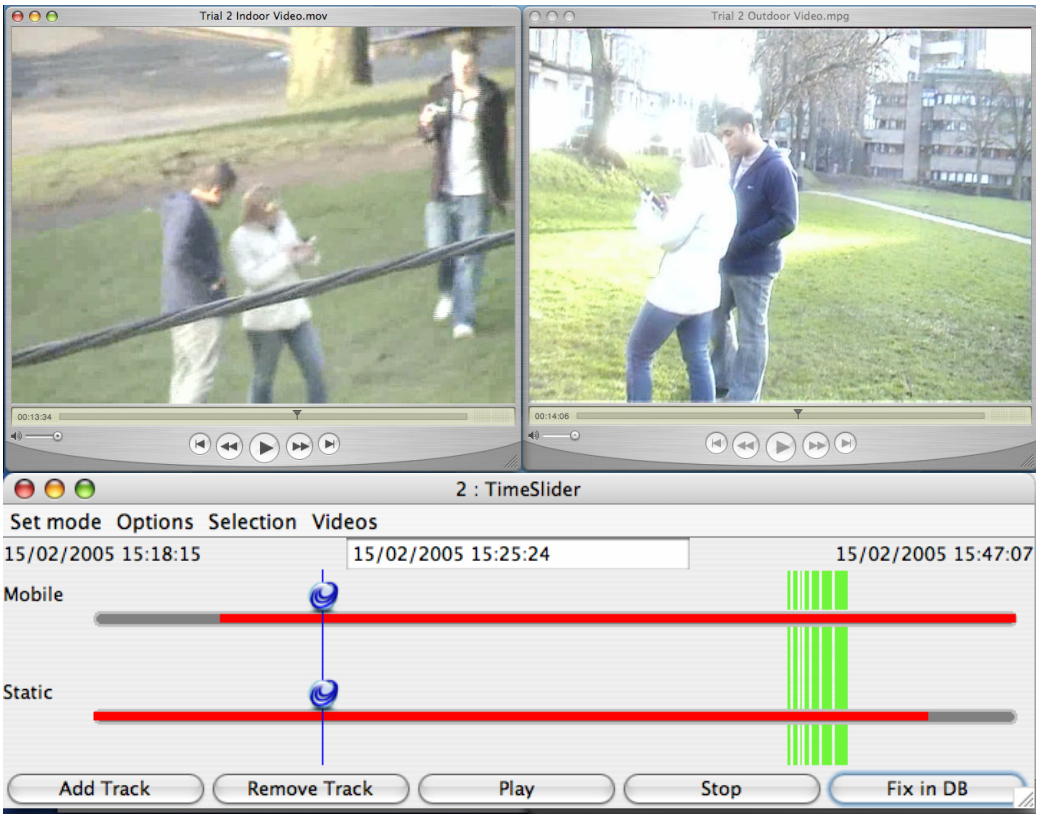

It is also from this component that new media clips are added to the database. To add a video it must first be loaded into QuickTime then synchronised to the data. This is done by entering into the textfield at the top of the component the 'real world' time at which the footage was recorded. This is often most simply achieved by recording a clock at some point and then skipping to that point in the footage to synchronise. Doing this by recording time shown on a trial device as in the figure ensures that the video is synchronised to that device.

Google Earth Link

The component requires Google Earth (GE) to be installed on the user's machine. It is available for free download at http://earth.google.com. Replayer's Google Earth Link component works by creating keyhole markup language (KML) files, then serving them to GE over HTTP. To run the component, launch it from the Control Unit and then open GE. In GE, open the *Add* menu and choose *Network Link*. Give it the name Replayer and set its location to "http://127.0.0.1:8081/Replayer". The network link should appear in GE's places menu; select its checkbox to view the data on the map.

The data table sent to the GE component must contain positional columns, and the X-Coordinate and Y-Coordinate combo-boxes should be set to specify which database columns contain the positional information. These can be lat-long values or OSGB references, with a command in the *Options* menu allowing the user to set the component for the type of data provided. Another menu command sets the component to draw cumulative trails between points in playback mode, so as the playback progresses, lines for each participant get longer as they show the routes taken during the recorded trial.

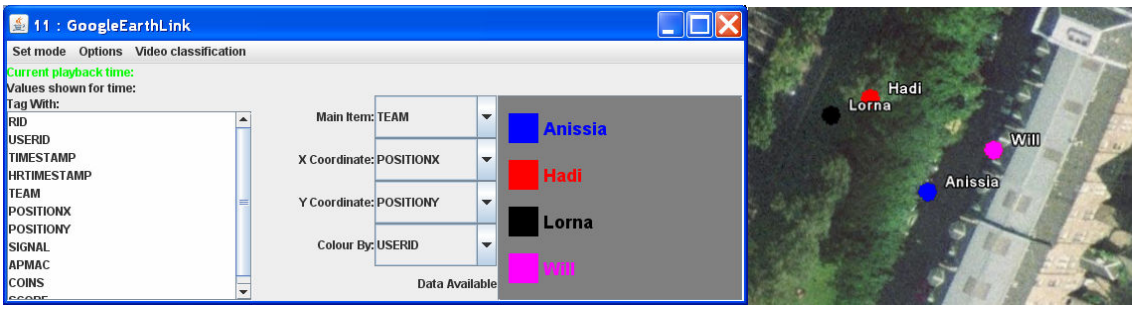

In Replayer's playback mode, a dot is shown on the map for each player, and these update in synchronisation with the video playback to show the players' positions as they move around the map. In overview mode, a dot is drawn for every row of the table in the current selection. This allows, for example, a single user's entire trail to be observed, or the location of every occurrence of a particular event to be displayed. Points are coloured by participant name. Additional dimensions of data can be tagged onto the points by making use of the description bubbles provided for each marker on the map.

Playback Controls Component

It may be the case that a dataset does not contain video, but it is still desirable to view the data in 'playback mode' to see how state and map data progressed during a trial. The Playback Controls component allows this. The user specifies a start time and playback speed and the component sends out a series of timestamps to the visualisation components, instructing them to update to show the state at that moment.

Data can be pushed to this component in the normal way, with it initialising its 'start time' field to the timestamp of the earliest value in the table it receives. Pulse times can be set, to alter the frequency with which the component sends its 'update' message to the other visualisations.